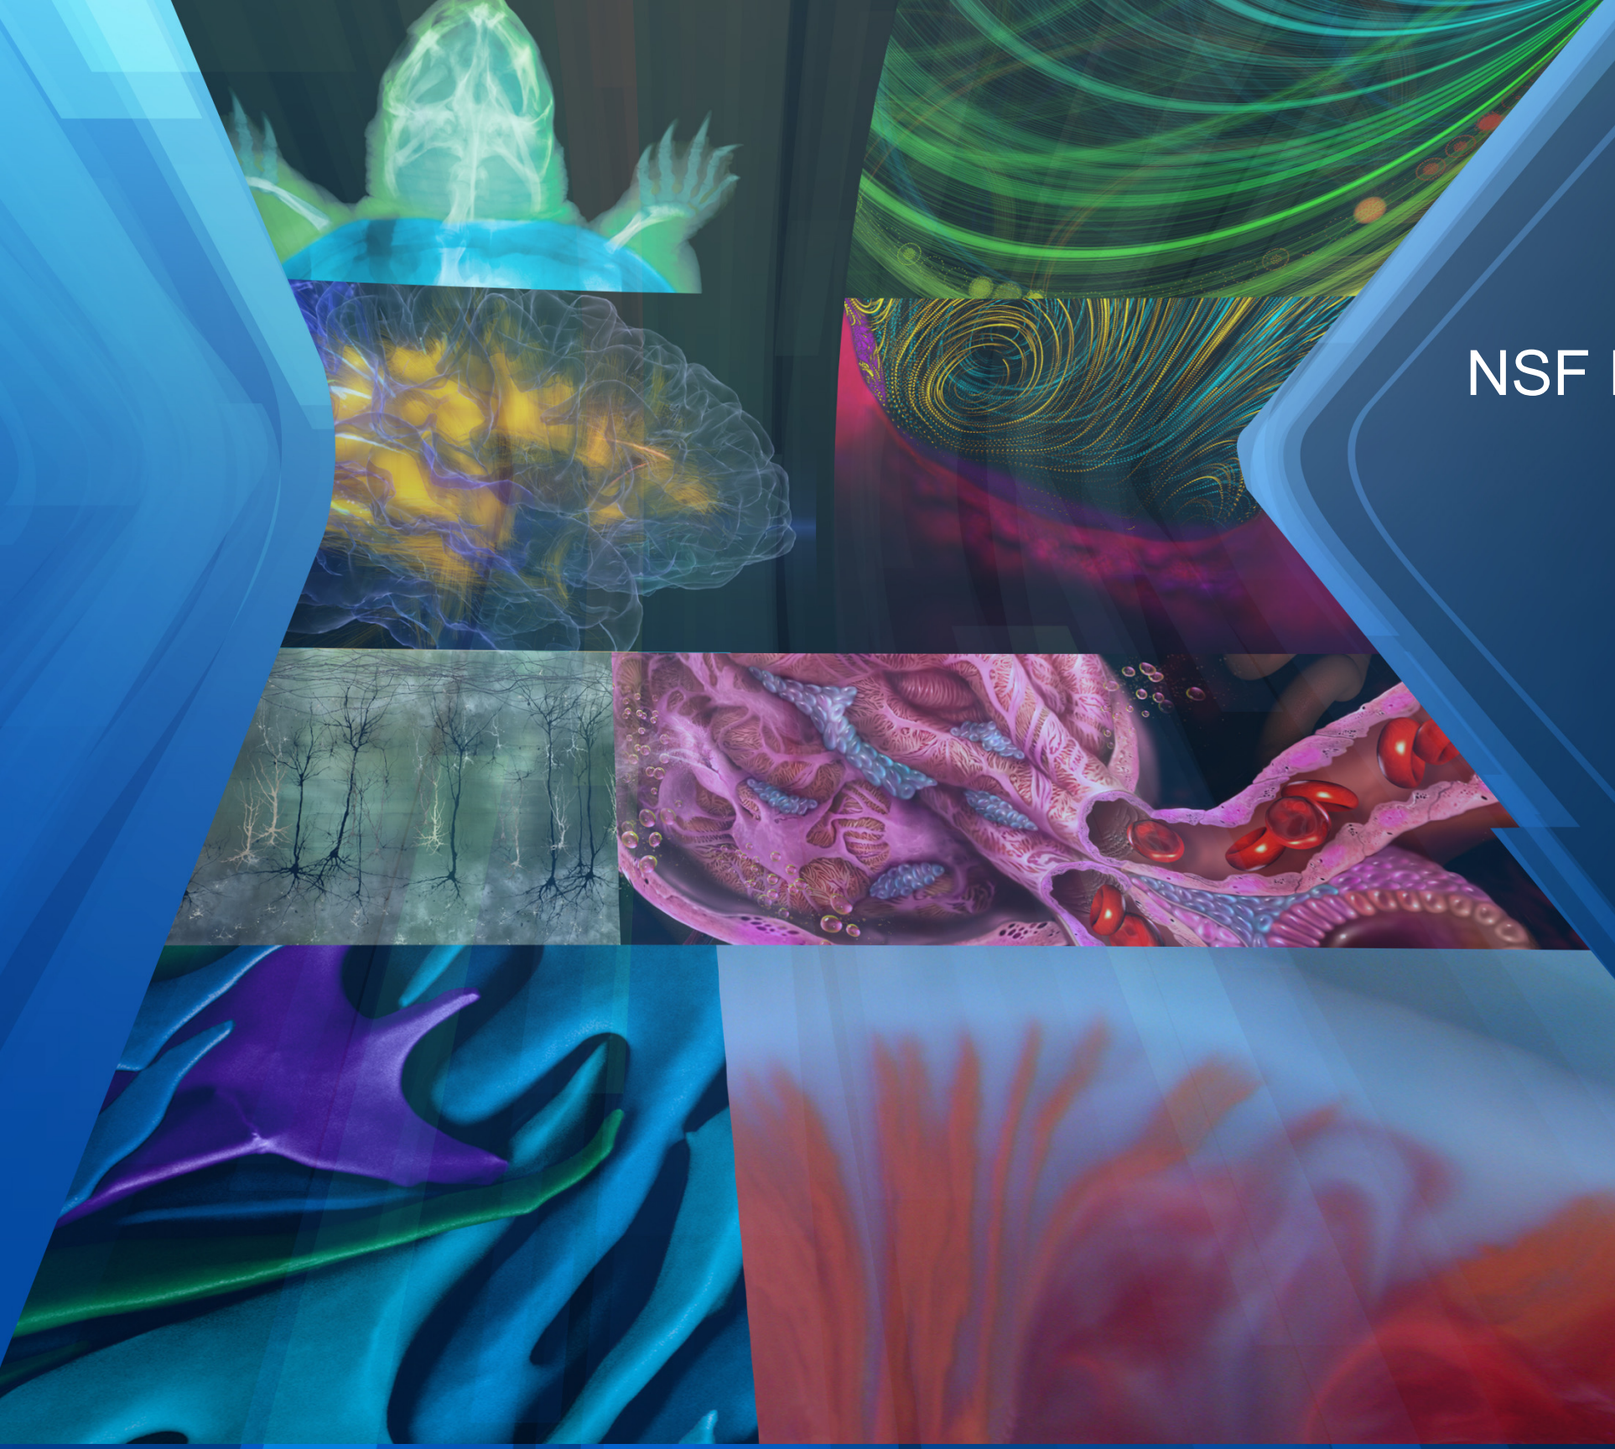

NSF DCL 21-108: Towards an Equitable National Cyberinfrastructure

Webinar Session
August 27, 2021, 2pm ET

NSF DCL 21-108
Towards an Equitable National Cyberinfrastructure

Agenda

- Team Introductions
- Background and Mission
- EPSCoR program
- DCL 21-108 Goals
- CC* 21-158 Program
- The Quilt
- Q&A Session

https://www.nsf.gov/pubs/2021/nsf21108/nsf21108.jsp?org=NSF

NSF Mission

"To promote the progress of science; to advance the national health, prosperity, and welfare; to secure the national defense..."

NSF by the Numbers (2021)

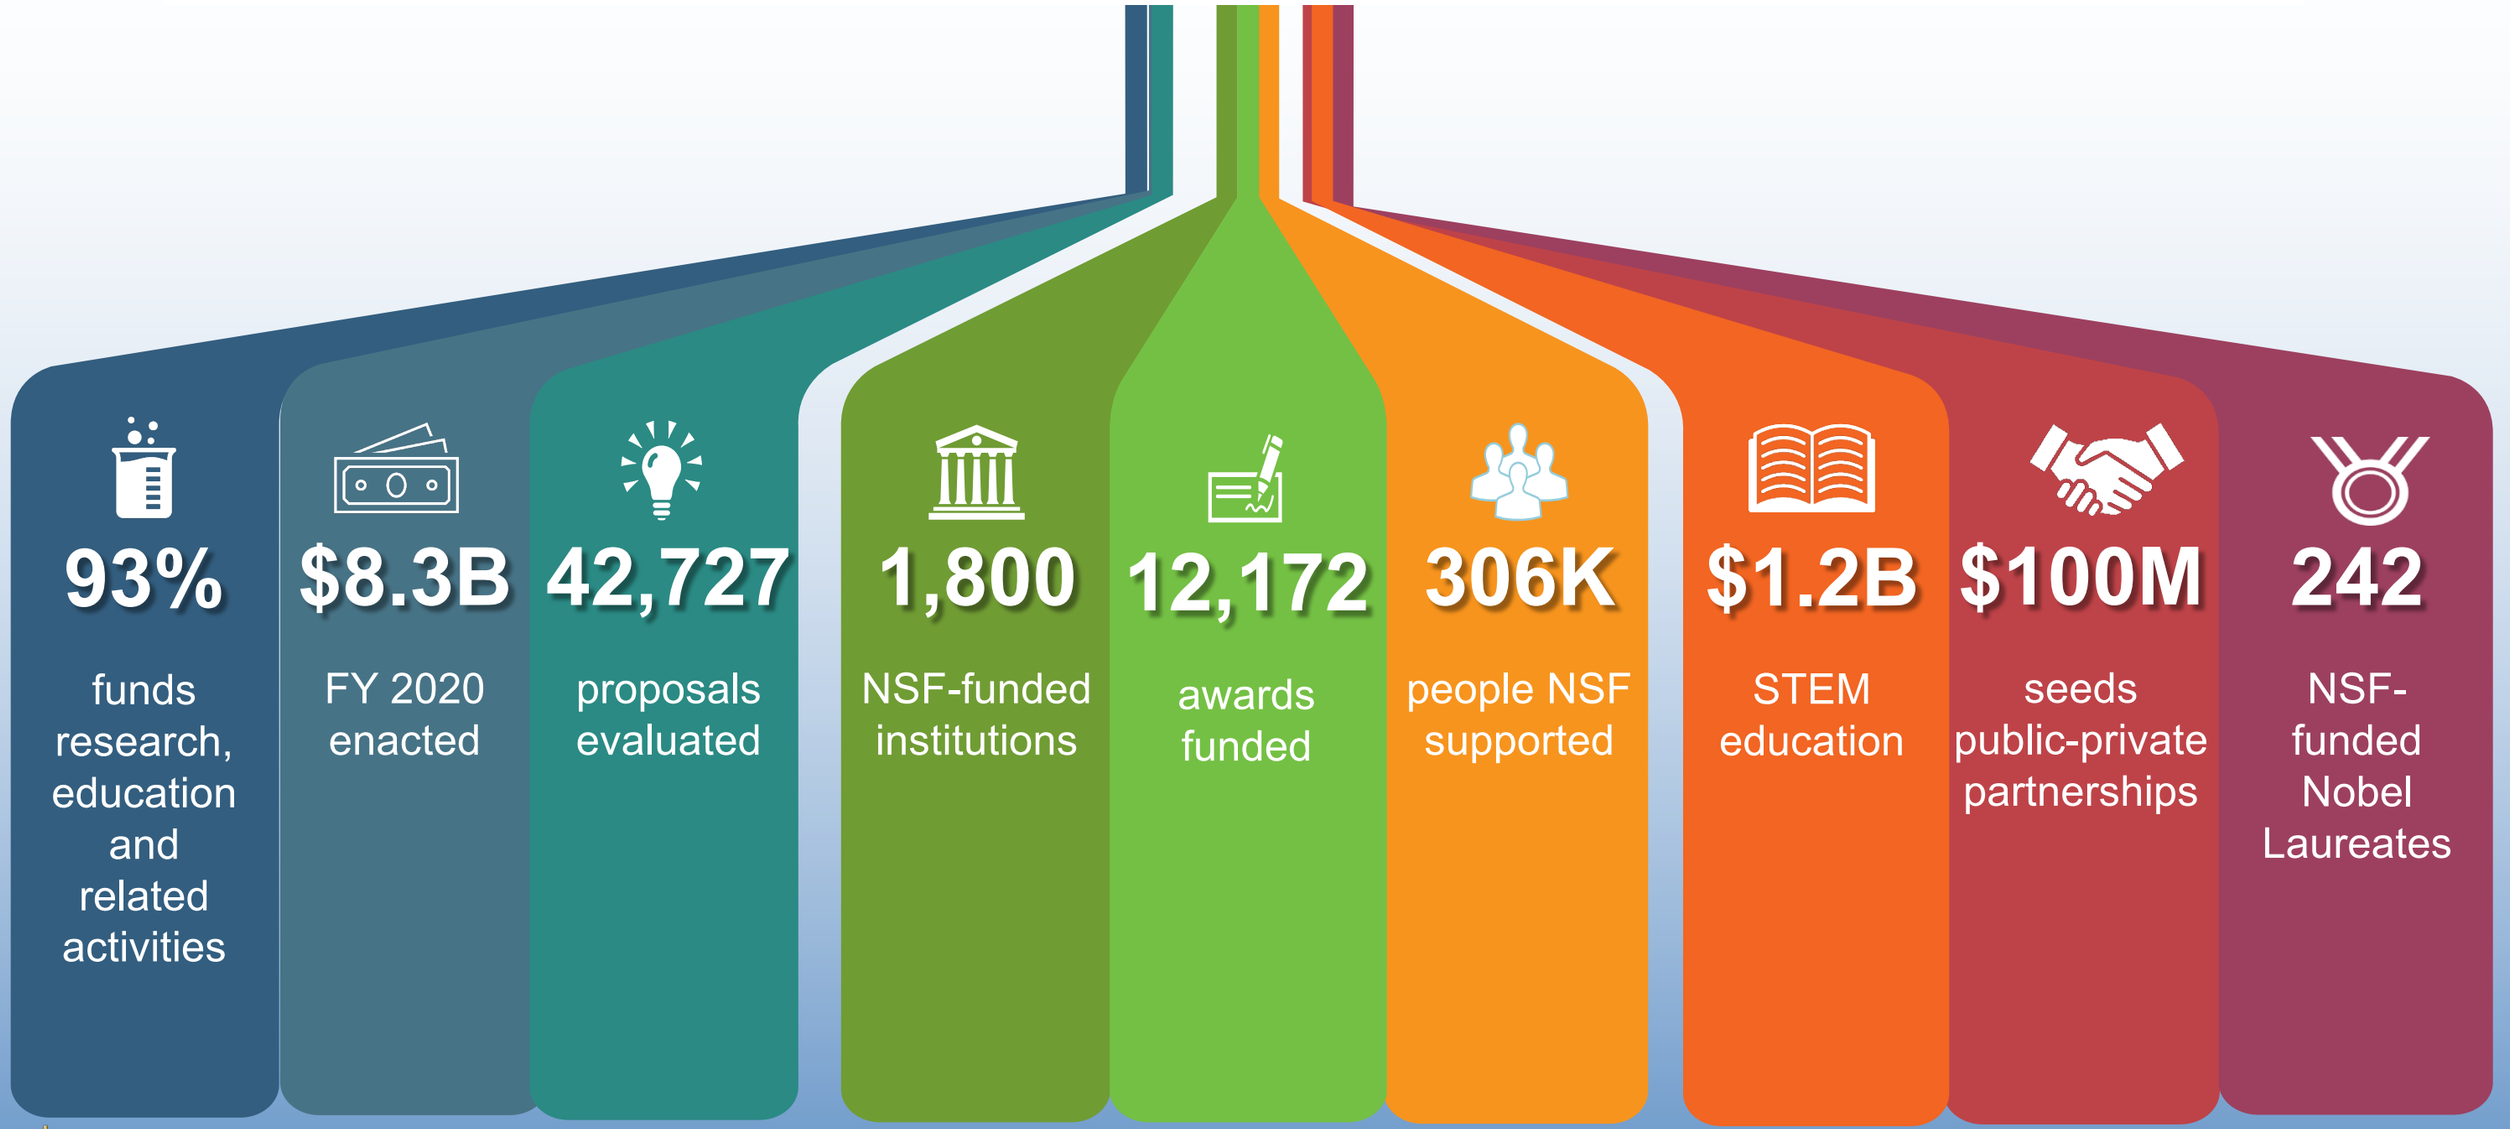

NSF Director's Strategic Vision

Dr. Sethuraman Panchanathan, the 15th director of NSF, has a strategic vision that includes three pillars rooted in the legacy of NSF's founding

- **Global Competition:** maintaining global leadership in the scientific enterprise
- **Missing Millions:** ensuring accessibility and inclusivity for aspiring scientists and engineers in STEM fields and (demographic, geographical and socio-economic)
  - Scale existing pathways and create new pathways to science and engineering
- **Industries of the Future:** advancing the frontiers of research into the future

Foundation for theses pillars are on Partnerships and Innovation

EPSCoR Mission

Enhances research competitiveness of targeted jurisdictions (states, territories, commonwealth) by strengthening STEM capacity and capability

Goals

- ***Catalyze research*** capability across and among jurisdictions
- ***Establish STEM*** professional development pathways
- ***Broaden participation*** of diverse groups and institutions in STEM
- ***Effect engagement*** in STEM at national and global levels
- ***Impact*** jurisdictional economic development

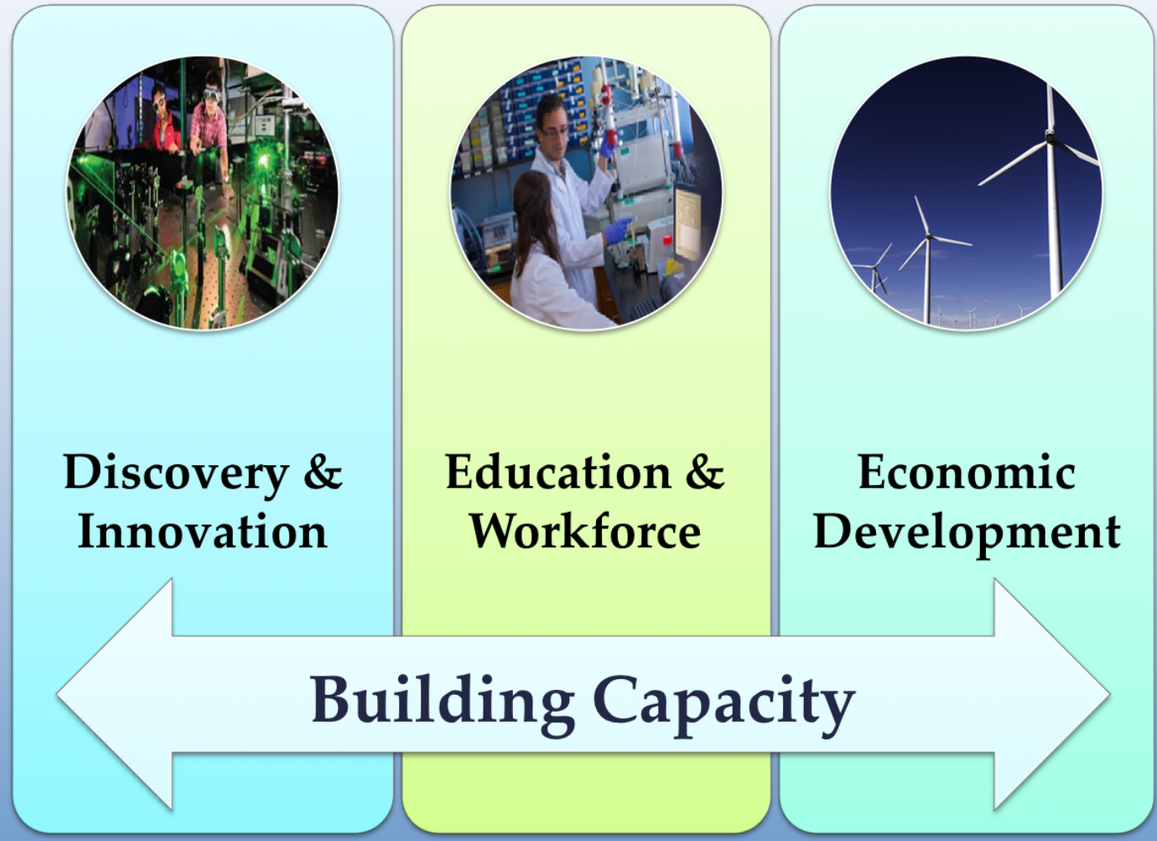

https://www.nsf.gov/od/oia/programs/epscor/index.jsp

EPSCoR Jurisdictions

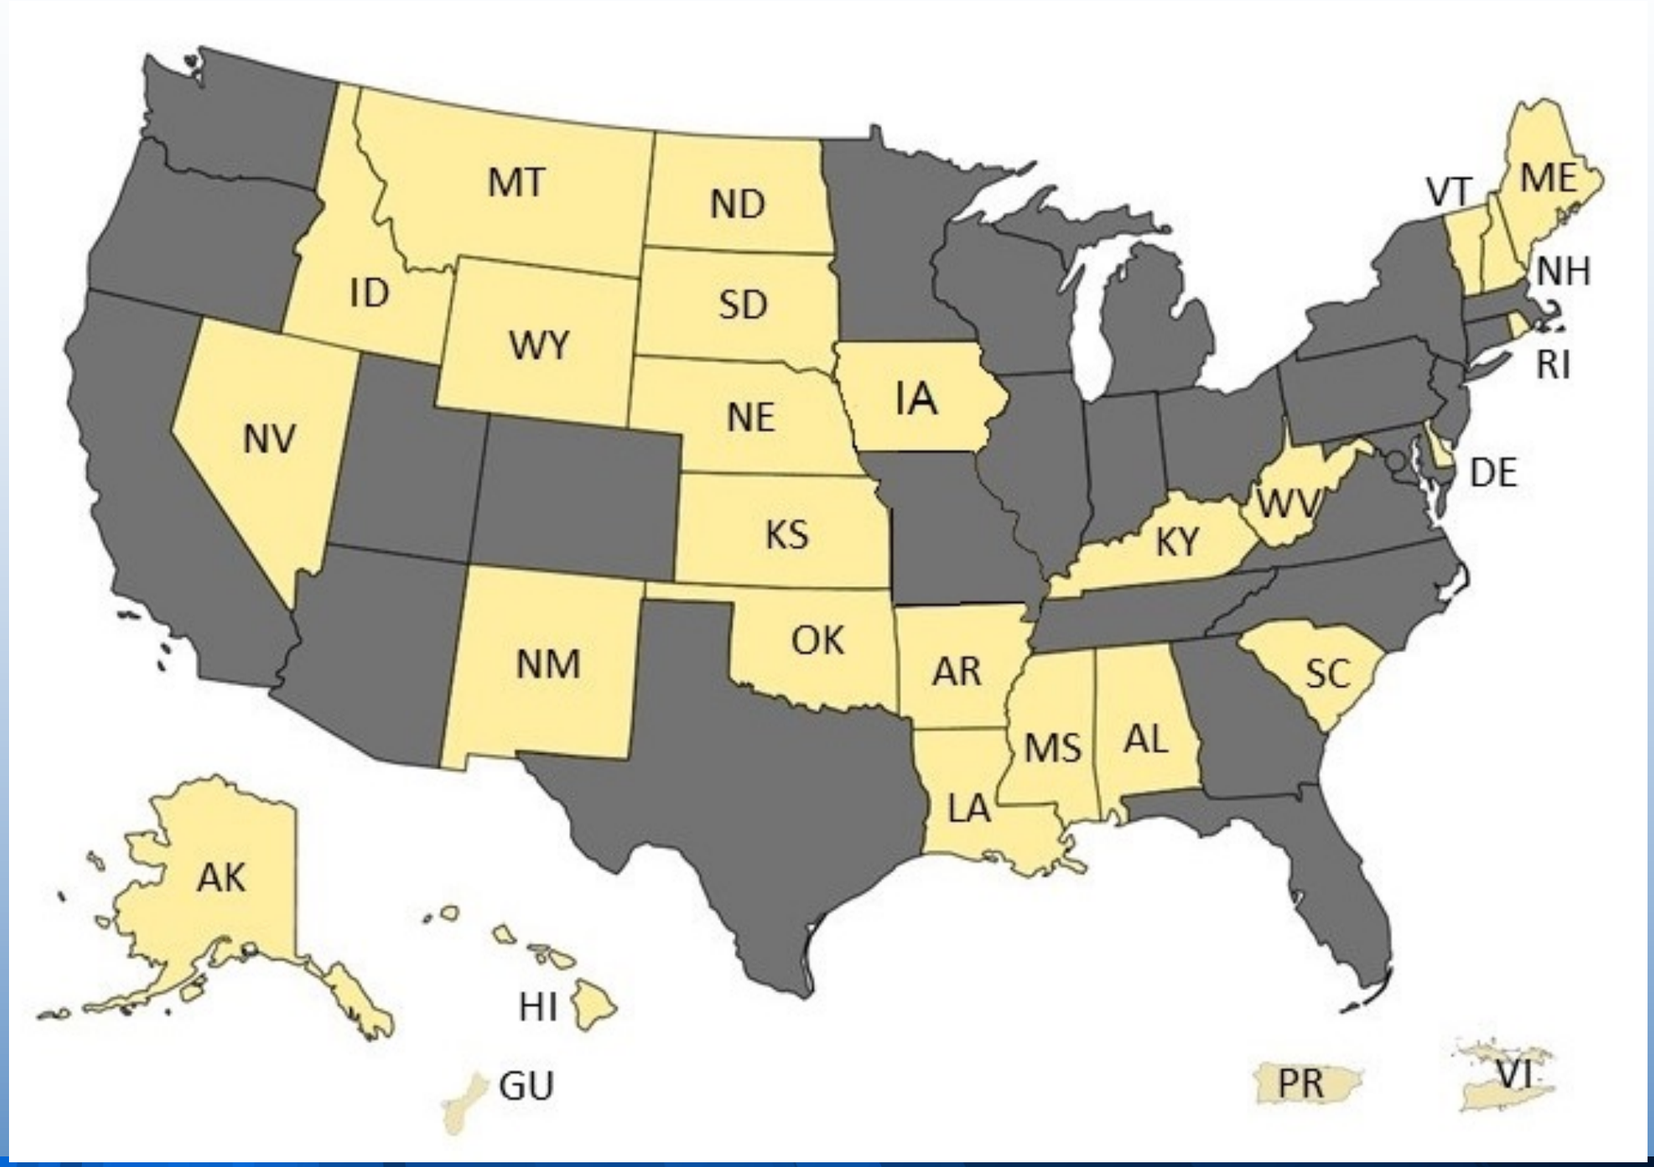

https://www.nsf.gov/od/oia/programs/epscor/Eligibility_Tables/FY2022_Eligibility.pdf

EPSCoR Investment Strategies

- **Research Infrastructure Improvement (RII)**

  Support physical, human, and cyber infrastructure within academic institutions across each jurisdiction

- **Co-Funding w/ NSF Directorates & Offices**

  Meritorious proposals reviewed in other NSF programs that also satisfy EPSCoR programmatic criteria

- **Outreach and Workshops**

  Interaction among EPSCoR Community and NSF to build mutual awareness and develop areas of potential strength

EPSCoR Focus on Collaborative Partnerships

RII Track-1

- Collaborations among regional and national EPSCoR jurisdiction-based organizations are encouraged, as are partnerships with nationally recognized centers of R&D activity, such as federal and industrial R&D laboratories, NSF-sponsored research centers, and academic institutions with nationally-recognized research capabilities.

RII Track-2

- Builds nationally and internationally competitive collaborative teams of EPSCoR investigators by providing a mechanism to coalesce investigator expertise into a critical mass for a sustained, effective research and education partnership.

RII Track-4

- Provides opportunities for non-tenured investigators to further develop their individual research potential through extended collaborative visits to the nation's premier private, governmental, or academic research centers. During these visits, the EPSCoR Research Fellows learn new techniques, develop new collaborations and advance existing partnerships.

Co-Funding w/ NSF Directorates

- Joint support of meritorious proposals from EPSCoR Institutions (New PI; Collaborative/Multidisciplinary; Synergy with NSF Priorities; Broadening Participation; Instrumentation to Increase Research Capacity; K-12 Student & Teacher Training; Integration of Research & Education).

Workshops and Outreach

- Support EPSCoR community-wide activities designed to explore opportunities in emerging areas of STEM research; Share best practices in planning and implementation in areas of importance to EPSCoR jurisdictions. Supports NSF staff travel for collaborations and partnerships with the EPSCoR community and builds mutual awareness

EPSCoR Funding ($M)

Activity	FY16	FY17	FY18	FY19	FY20
RII	130.40	135.80	142.20	144.90	148.60
Track-1	68.92	83.34	108.66	99.08	93.86
Track-2	56.63	43.51	23.70	42.66	42.31
Track-4	0	5.90	6.70	0	9.02
Other	4.85	3.05	3.14	3.16	3.41
Co-funding	28.5	24.9	27.6	30.60	41.9
Outreach & Workshops	1.1	2.1	0.8	0.1	1.2
Total*	160.0	162.8	170.6	175.6	191.7

* May not add due to rounding

Increasing Access to Cyberinfrastructure
NSF and EPSCoR Priority

Invest time to understand the needs of the Researchers, Educators, and IT staff at the institution-level.

- Develop institutional partnerships to capitalize on resources that are already in-place and work together to acquire those that are not.
- Learn from community experiences of the challenges and obstacles faced by the end-source users (e.g. researchers, educators, students, IT administrative staff, ...).
- Identify opportunities for in-depth immersion to truly understand the campus' computing resources and environments and their unique needs.

IT and computing resources are not one-size fits all.

DCL 21-108: Towards an Equitable National Cyberinfrastructure

- Partnership effort NSF effort between the **Campus Cyberinfrastructure (CC*) Program** (Office of Advanced Cyberinfrastructure (OAC)) and the **Established Program to Stimulate Competitive Research (EPSCoR)** (Office of Integrative Activities (OIA))
- Response to targeted activities under the **American Rescue Plan Act, 2021**
  - Targeted specifically to those that have been **disproportionately** impacted by the pandemic
- DCL focus on **Cyberinfrastructure support efforts** for institution and individuals disproportionately affected by the pandemic
  - Examples: enhance cyber connectivity, mitigating gaps in access to cyber resources, etc.

Goal: Invest in coordinated campus-level networking and cyberinfrastructure improvements, innovation, integration, and engineering for science and engineering applications and distributed research projects.

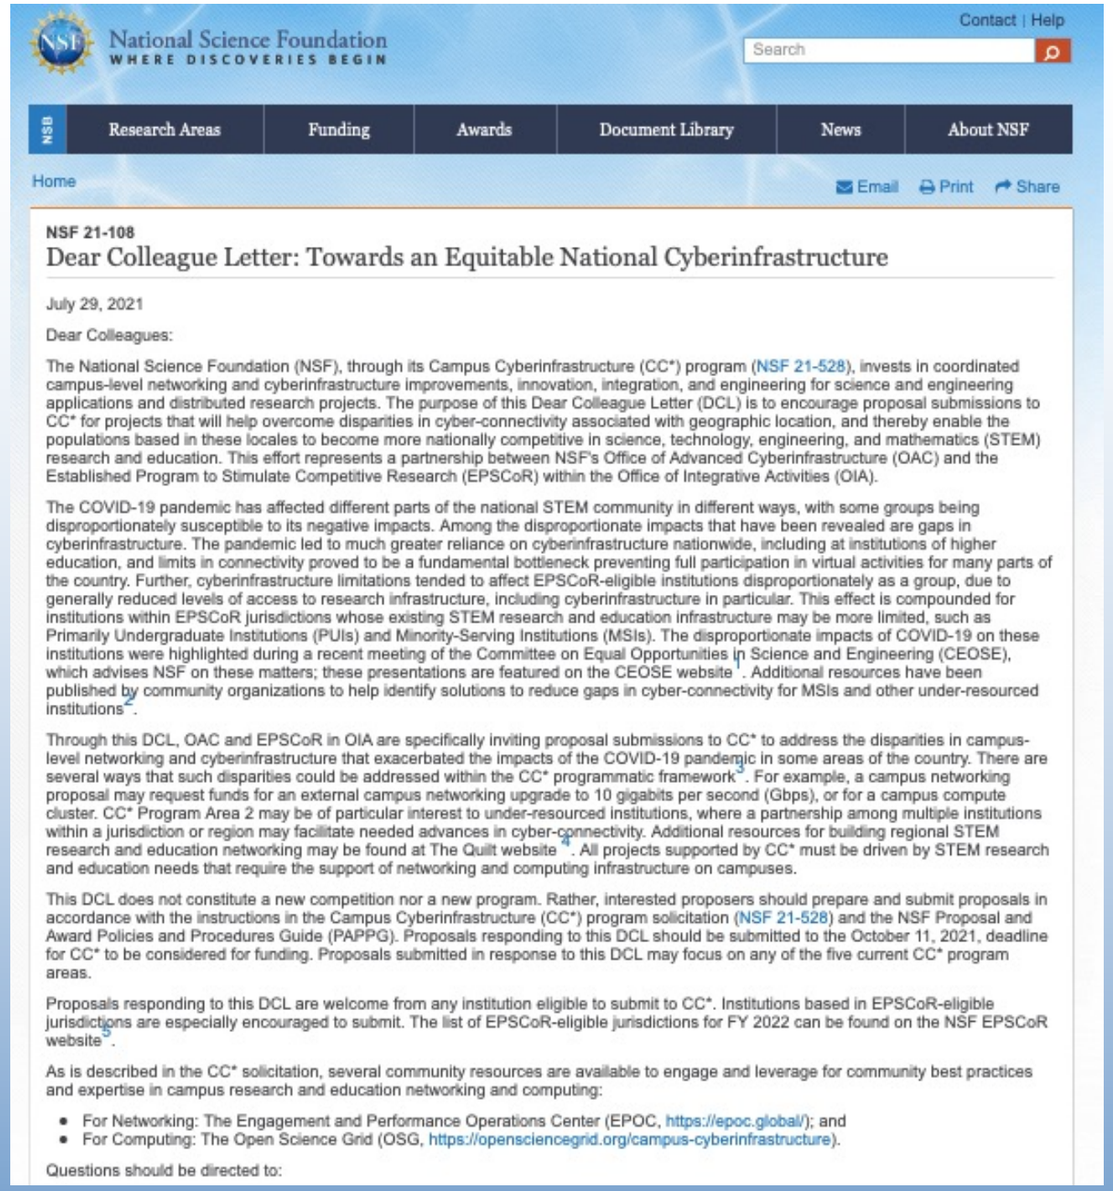

National Science Foundation
WHERE DISCOVERIES BEGIN

Contact | Help

Search

NSB | Research Areas | Funding | Awards | Document Library | News | About NSF

Home

Email Print Share

NSF 21-108

Dear Colleague Letter: Towards an Equitable National Cyberinfrastructure

July 29, 2021

Dear Colleagues:

The National Science Foundation (NSF), through its Campus Cyberinfrastructure (CC*) program (NSF 21-528), invests in coordinated campus-level networking and cyberinfrastructure improvements, innovation, integration, and engineering for science and engineering applications and distributed research projects. The purpose of this Dear Colleague Letter (DCL) is to encourage proposal submissions to CC* for projects that will help overcome disparities in cyber-connectivity associated with geographic location, and thereby enable the populations based in these locales to become more nationally competitive in science, technology, engineering, and mathematics (STEM) research and education. This effort represents a partnership between NSF's Office of Advanced Cyberinfrastructure (OAC) and the Established Program to Stimulate Competitive Research (EPSCoR) within the Office of Integrative Activities (OIA).

The COVID-19 pandemic has affected different parts of the national STEM community in different ways, with some groups being disproportionately susceptible to its negative impacts. Among the disproportionate impacts that have been revealed are gaps in cyberinfrastructure. The pandemic led to much greater reliance on cyberinfrastructure nationwide, including at institutions of higher education, and limits in connectivity proved to be a fundamental bottleneck preventing full participation in virtual activities for many parts of the country. Further, cyberinfrastructure limitations tended to affect EPSCoR-eligible institutions disproportionately as a group, due to generally reduced levels of access to research infrastructure, including cyberinfrastructure in particular. This effect is compounded for institutions within EPSCoR jurisdictions whose existing STEM research and education infrastructure may be more limited, such as Primarily Undergraduate Institutions (PUIs) and Minority-Serving Institutions (MSIs). The disproportionate impacts of COVID-19 on these institutions were highlighted during a recent meeting of the Committee on Equal Opportunities in Science and Engineering (CEOSE), which advises NSF on these matters; these presentations are featured on the CEOSE website[1]. Additional resources have been published by community organizations to help identify solutions to reduce gaps in cyber-connectivity for MSIs and other under-resourced institutions[2].

Through this DCL, OAC and EPSCoR in OIA are specifically inviting proposal submissions to CC* to address the disparities in campus-level networking and cyberinfrastructure that exacerbated the impacts of the COVID-19 pandemic in some areas of the country. There are several ways that such disparities could be addressed within the CC* programmatic framework[3]. For example, a campus networking proposal may request funds for an external campus networking upgrade to 10 gigabits per second (Gbps), or for a campus compute cluster. CC* Program Area 2 may be of particular interest to under-resourced institutions, where a partnership among multiple institutions within a jurisdiction or region may facilitate needed advances in cyber-connectivity. Additional resources for building regional STEM research and education networking may be found at The Quilt website[4]. All projects supported by CC* must be driven by STEM research and education needs that require the support of networking and computing infrastructure on campuses.

This DCL does not constitute a new competition nor a new program. Rather, interested proposers should prepare and submit proposals in accordance with the instructions in the Campus Cyberinfrastructure (CC*) program solicitation (NSF 21-528) and the NSF Proposal and Award Policies and Procedures Guide (PAPPG). Proposals responding to this DCL should be submitted to the October 11, 2021, deadline for CC* to be considered for funding. Proposals submitted in response to this DCL may focus on any of the five current CC* program areas.

Proposals responding to this DCL are welcome from any institution eligible to submit to CC*. Institutions based in EPSCoR-eligible jurisdictions are especially encouraged to submit. The list of EPSCoR-eligible jurisdictions for FY 2022 can be found on the NSF EPSCoR website[5].

As is described in the CC* solicitation, several community resources are available to engage and leverage for community best practices and expertise in campus research and education networking and computing:

- For Networking: The Engagement and Performance Operations Center (EPOC, https://epoc.global/); and
- For Computing: The Open Science Grid (OSG, https://opensciencegrid.org/campus-cyberinfrastructure).

Questions should be directed to:

NOTE: Proposals responding to this DCL are welcome from any institution eligible to submit to CC. Institutions based in EPSCoR-eligible jurisdictions are especially encouraged to submit.*

NSF DCL 21-108 Goals

(collaborative effort from the EPSCoR and Campus Cyberinfrastructure (CC*) Programs)

Respond to the NSF DCL 21-108 on establishing an equitable national Cyberinfrastructure

Foster innovative collaborations to address disparities and enhance regional and national Cyberinfrastructure goals

Explore topics in emerging areas of science and engineering to enable equitable partnerships

Bring together communities of thought to explore on recent research or education findings

Expose researchers and trainees at underserved institutions to new collaborations, tools & techniques

NSF CISE by the numbers, FY 2020

25% success rate

$1,011 M FY 2020 enacted budget

7,932 proposals

432 Panels

1,971 awards

19,699 people supported

369 institutions supported

8,164 senior researchers

1,134 other professionals

561 postdoctoral associates

6,622 graduate students

3,218 undergraduate students

NSF funds > 85% of federally-funded academic CS research in the US.

(*Source: NCSES*)

ACCI Task Force on Campus Bridging

- *Strategic Recommendation to the NSF #3: The National Science Foundation should create a new program funding high-speed (currently 10 Gbps) connections from campuses to the nearest landing point for a national network backbone. The design of these connections must include support for dynamic network provisioning services and must be engineered to support rapid movement of large scientific data sets."* - pg. 6, National Science Foundation Advisory Committee for Cyberinfrastructure Task Force on Campus Bridging, Final Report, March 2011
- www.nsf.gov/od/oci/taskforces/TaskForceReport_CampusBridging.pdf
- Also see Campus Bridging Technologies Workshop: Data and Networking Issues Workshop Report. G.T. Almes, D. Jent and C.A. Stewart, eds., 2011, http://hdl.handle.net/2022/13200

Building the research network substrate – campus level: CC*
~350 Awards across 49 states/jurisdictions 2012-today

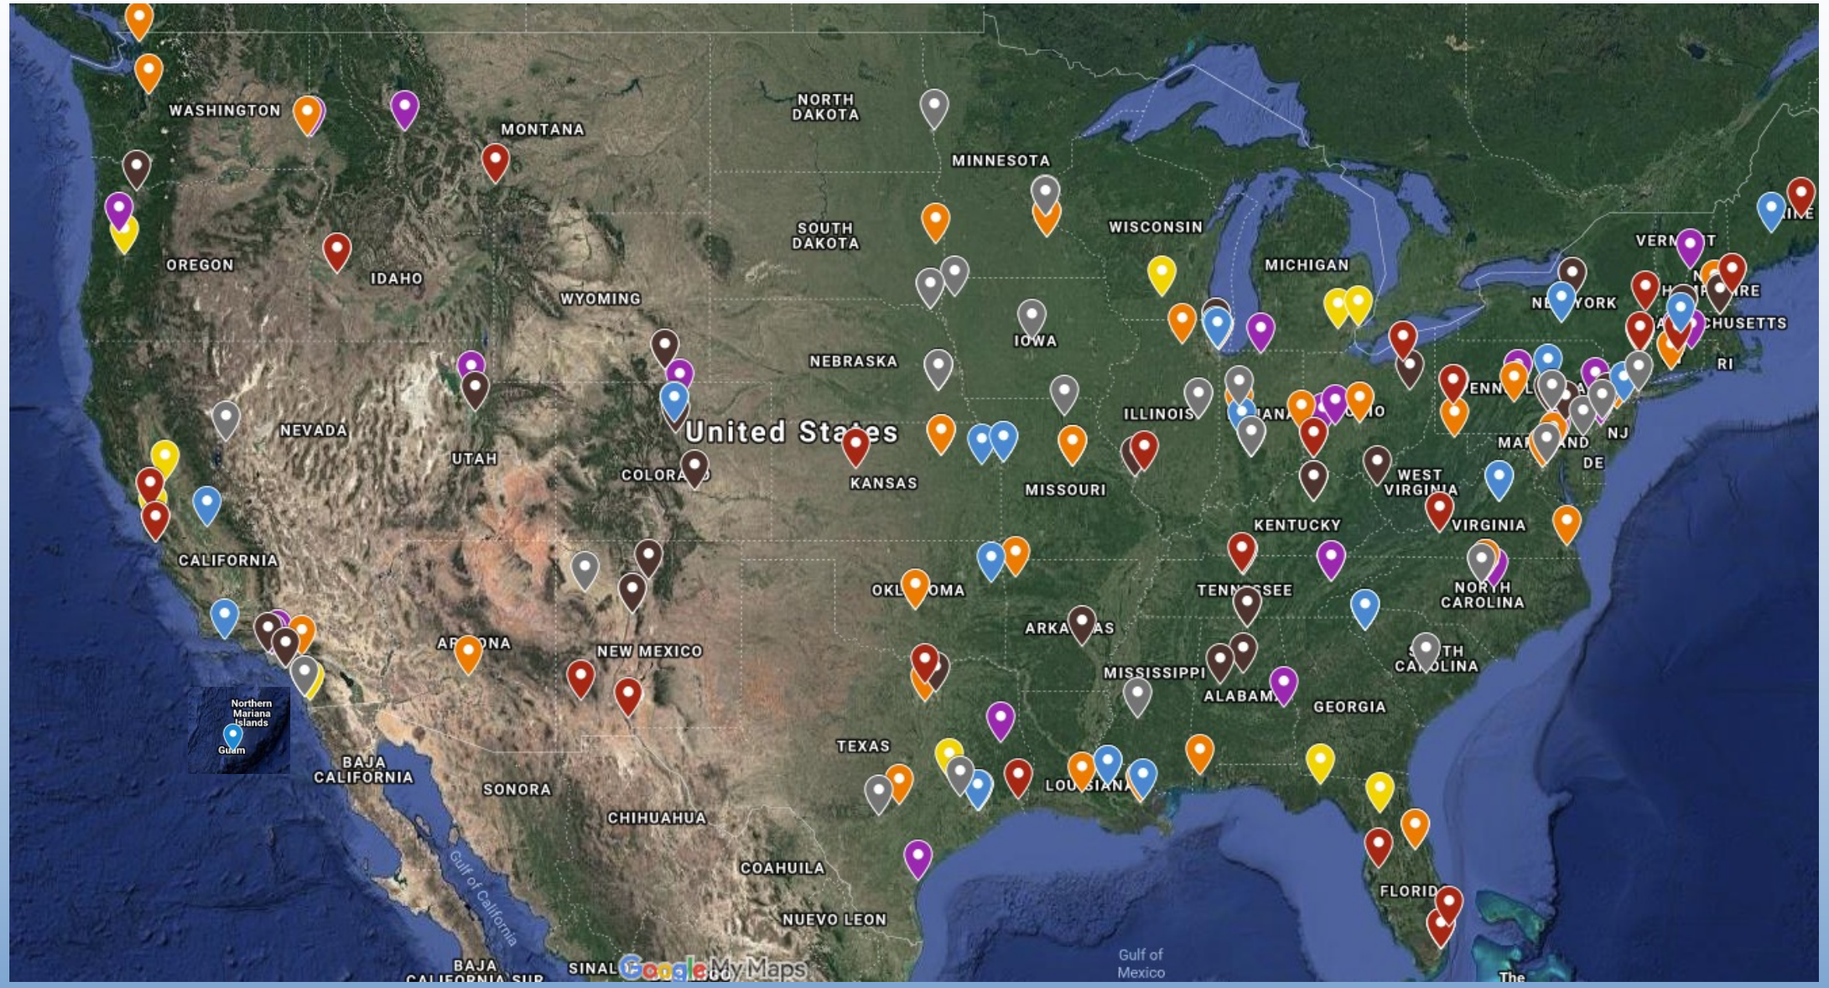

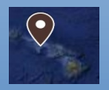

CAMPUS CYBERINFRASTRUCTURE (CC*): Upgrades networking capabilities for >200 campuses across 45 states to support science applications and distributed research projects

Campus Cyberinfrastructure (CC*) NSF 21-528

The Campus Cyberinfrastructure (CC*) program invests in coordinated campus-level networking and cyberinfrastructure improvements, innovation, integration, and engineering for science applications and distributed research projects. Science-driven requirements are the primary motivation for any proposed activity.

Summary #s for NSF's Campus CI Program 2012-2020

- ~350 awards (not including workshops, EAGER)
- $110M+ invested over 9 years (FY12-FY21)
- Award categories [aggregate (FY20)]:
  - Campus Networking Upgrades: 143 (6)
  - Network Integration/Innovation: 60 (10)
  - Regional/Network Design: 37 (3)
  - CyberTeam / CI Engineer: 33 (5)
  - Compute: 33 (19)
  - Other: 35 (7 planning grants)

CC* 21-528 - Campus Cyberinfrastructure

- https://www.nsf.gov/pubs/2021/nsf21528/nsf21528.htm
- $13M-$15M in expected award funding
- **Proposals due ~~March 1 and~~ October 11, 2021**
- Area #1 – **Campus Network upgrades**
  - 10/100Gbps+ inter- and intra-campus networking
  - Re-design of campus border to prioritize science flows
  - Awards up to $500,000
- Area #2 – **Regional coordination for Small Institutions**
  - Establishing r&e network connectivity for multiple under-resourced institutions
  - Awards up to $1,000,000
- Area #3 – **Networking Integration and Applied Innovation**
  - Applied R&D in networking motivated by science use cases
  - Awards up to $500,000 (small) and $1,000,000 (large)
- Area #4 – **Campus Computing**
  - Shared cluster cycles for campus-wide science
  - Awards up to $400,000
- Area #5 – **Planning Grants and CI-Research Alignment**
  - Awards up to $100,000 (planning) and $200,000

Program-wide Criteria for CC* proposals

- **Science-driven requirements** are the primary motivation for any proposed activity. Proposals will be evaluated on the strength of the science enabled (including research and education) as drivers for investment and innovation in data networking infrastructure, innovation, and engineering.
- **A partnership among researchers/educators and campus IT leadership**
- All proposals submitted to the CC* program, with the exception of area (5), must include a **Campus CI plan** within which the proposed CI improvements are conceived, designed, and implemented in the context of a coherent campus-wide strategy and approach to CI that is integrated horizontally intra- campus and vertically with regional and national CI investments and best practices. This Campus CI plan must be included as a Supplementary Document and is limited to no more than 5 pages.
- https://fasterdata.es.net/campusCIplanning/

CC* Area#1 - Data Driven Networking Infrastructure for the Campus and Researcher

- network infrastructure improvements at the campus level
- network improvements include:
  - Network upgrades within a campus network to support a wide range of science data flows…
  - re-architecting a campus network to support large science data flows, for example by designing and building a "science DMZ" (see http://fasterdata.es.net/science-dmz/ for more information on the "science DMZ" approach)
  - Network connection upgrade for the campus connection to a regional optical exchange or point-of-presence that connects to a state/regional/national network aggregation point prioritizing support for research and education
- In addressing networking equipment choices and configurations matched for high-performance R&E networking environments, proposals are encouraged to leverage objective community expertise and experience available from the NSF-funded EPOC project at: https://epoc.global/cc/.

https://fasterdata.es.net/science-dmz/

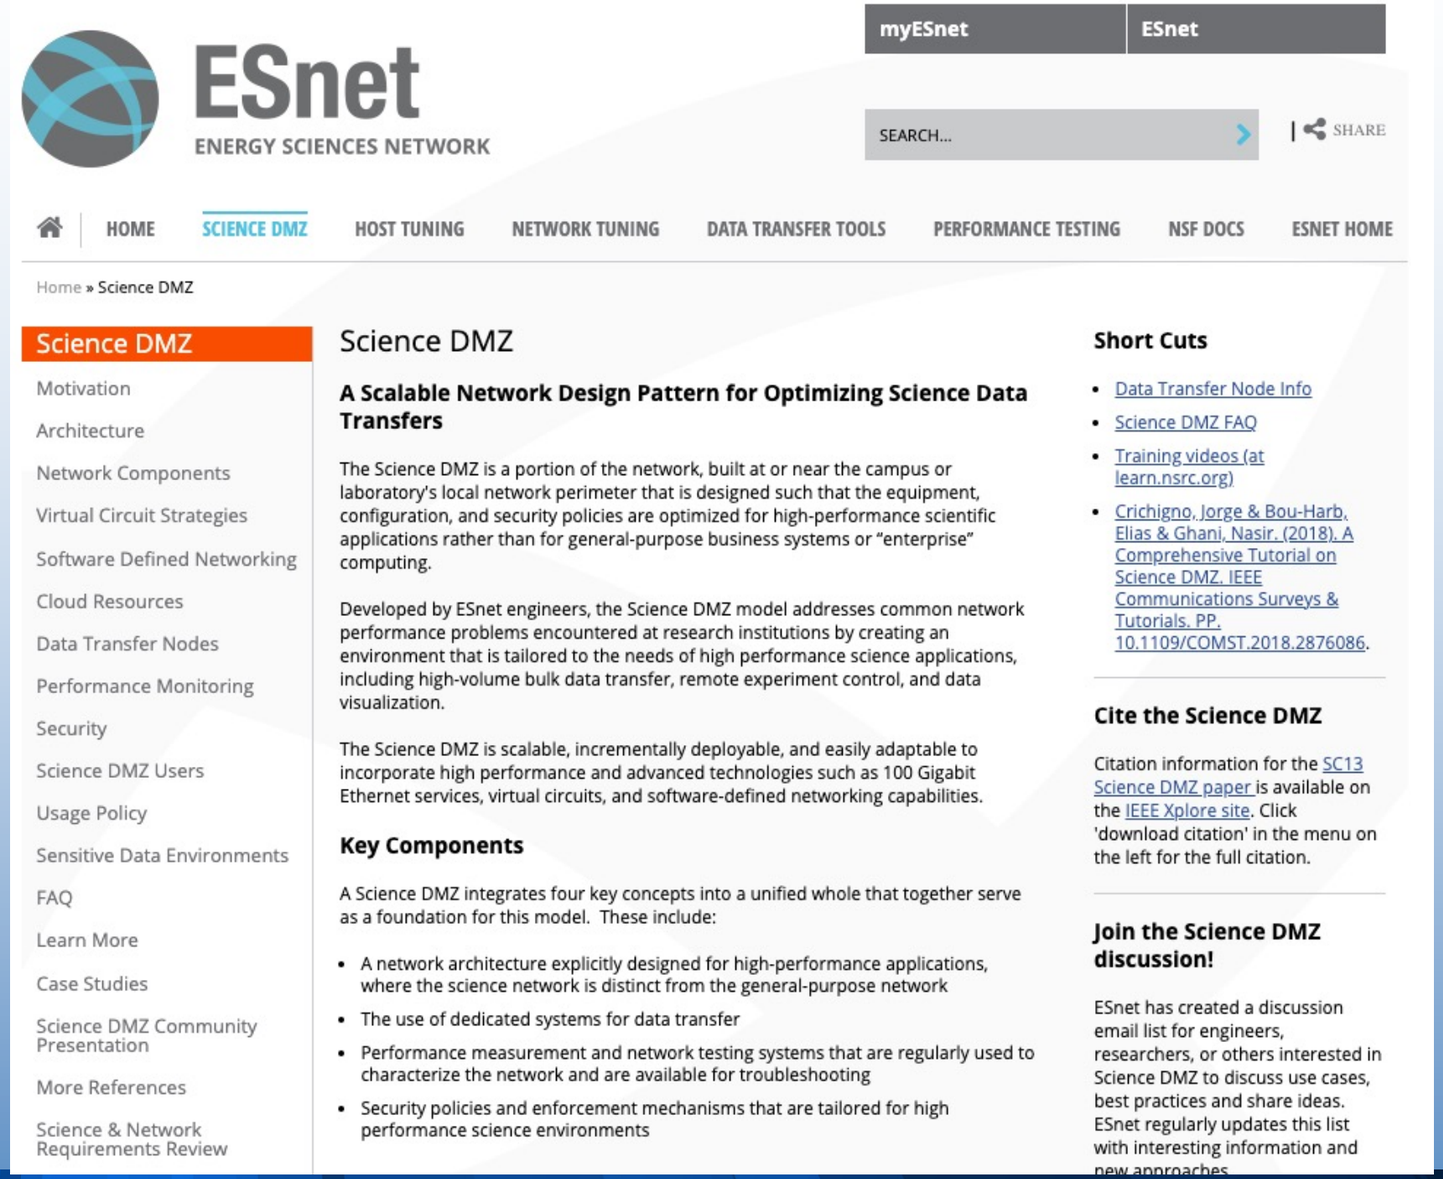

ESnet
ENERGY SCIENCES NETWORK

myESnet | ESnet

SEARCH... | SHARE

HOME | SCIENCE DMZ | HOST TUNING | NETWORK TUNING | DATA TRANSFER TOOLS | PERFORMANCE TESTING | NSF DOCS | ESNET HOME

Home » Science DMZ

Science DMZ
- Motivation
- Architecture
- Network Components
- Virtual Circuit Strategies
- Software Defined Networking
- Cloud Resources
- Data Transfer Nodes
- Performance Monitoring
- Security
- Science DMZ Users
- Usage Policy
- Sensitive Data Environments
- FAQ
- Learn More
- Case Studies
- Science DMZ Community Presentation
- More References
- Science & Network Requirements Review

Science DMZ

A Scalable Network Design Pattern for Optimizing Science Data Transfers

The Science DMZ is a portion of the network, built at or near the campus or laboratory's local network perimeter that is designed such that the equipment, configuration, and security policies are optimized for high-performance scientific applications rather than for general-purpose business systems or "enterprise" computing.

Developed by ESnet engineers, the Science DMZ model addresses common network performance problems encountered at research institutions by creating an environment that is tailored to the needs of high performance science applications, including high-volume bulk data transfer, remote experiment control, and data visualization.

The Science DMZ is scalable, incrementally deployable, and easily adaptable to incorporate high performance and advanced technologies such as 100 Gigabit Ethernet services, virtual circuits, and software-defined networking capabilities.

Key Components

A Science DMZ integrates four key concepts into a unified whole that together serve as a foundation for this model. These include:

- A network architecture explicitly designed for high-performance applications, where the science network is distinct from the general-purpose network
- The use of dedicated systems for data transfer
- Performance measurement and network testing systems that are regularly used to characterize the network and are available for troubleshooting
- Security policies and enforcement mechanisms that are tailored for high performance science environments

Short Cuts

- Data Transfer Node Info
- Science DMZ FAQ
- Training videos (at learn.nsrc.org)
- Crichigno, Jorge & Bou-Harb, Elias & Ghani, Nasir. (2018). A Comprehensive Tutorial on Science DMZ. IEEE Communications Surveys & Tutorials. PP. 10.1109/COMST.2018.2876086.

Cite the Science DMZ

Citation information for the SC13 Science DMZ paper is available on the IEEE Xplore site. Click 'download citation' in the menu on the left for the full citation.

Join the Science DMZ discussion!

ESnet has created a discussion email list for engineers, researchers, or others interested in Science DMZ to discuss use cases, best practices and share ideas. ESnet regularly updates this list with interesting information and new approaches

https://fasterdata.es.net/science-dmz/

Simple Science DMZ Diagram

A simple Science DMZ has several essential components. These include dedicated access to high-performance wide area networks and advanced services infrastructures, high-performance network equipment, and dedicated science resources such as Data Transfer Nodes. A notional diagram of a simple Science DMZ showing these components, along with data paths, is shown below:

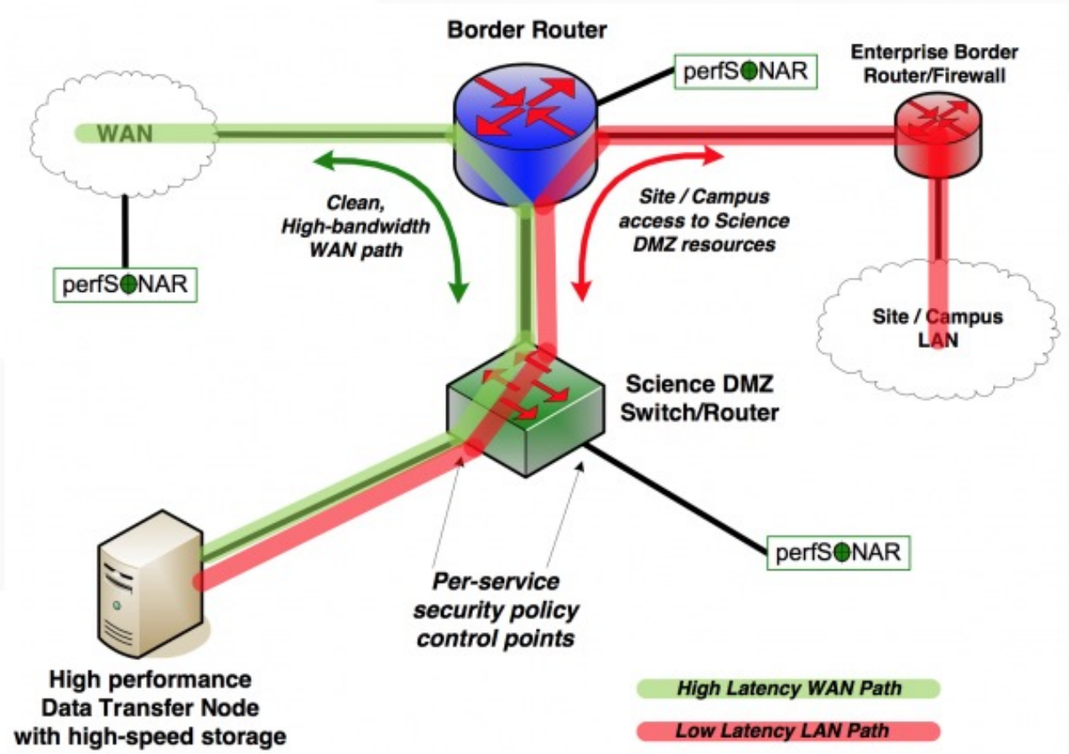

The essential components and a simple architecture for a Science DMZ are shown in the Figure above. The Data Transfer Node (DTN) is connected directly to a high-performance Science DMZ switch or router, which is connected directly to the border router. The DTN's job is to efficiently and effectively move science data to and from remote sites and facilities, and everything in the Science DMZ is aimed at this goal. The security policy enforcement for the DTN is done using access control lists on the Science DMZ switch or router, not on a separate firewall.

CC* Area#2 – Regional Connectivity for Small Institutions

- This area supports broadening participation and significantly widening the set of institutions connected to the regional and national research and education network fabric. **This area specifically targets groups of smaller institutions with fundamental challenges in networking infrastructure and resources**. This area supports increased research and education (R&E) network connectivity across smaller institutions **coordinated and led by a Regional Optical Network (RON) or a leadership institution in R&E networking in the region**.

CC* Area#2 – Regional Connectivity for Small Institutions of Higher Education

- This area solicits proposals led by established regional and state R&E data networks and data network-based consortia. Example entities are listed as members of the national regional network consortium called the Quilt (see https://www.thequilt.net/about-us/the-quilt-participants/).
- For areas of the US without a state or regional level coordinating entity and associated structure and network infrastructure, proposals will be accepted from self-declared leadership universities. An institution may also lead a proposal in regions with an established RON with documented coordination with the RON.
- **Proposals are required to address campus networking needs spanning multiple under-resourced institutions.** Proposals addressing a single institution are not allowed to submit to this area and will be returned without review. Proposals may choose to apply an alternative design framework to the conventional single institution context in Area (1) and consider an aggregation model where some or all associated resources and services (e.g., Science DMZ) are centralized at a regional level.

CC* Area#3 - Network Integration and Applied Innovation

- Goal - take advantage of research results, prototypes, and emerging innovations to use them to enable specified researchers in a networking context.
- Proposals in this area are expected to reflect innovation in advanced networking. As a result, this area is not appropriate for projects whose costs are dominated by equipment purchases.
- **Proposals in this area support the development and integration of innovative networking capabilities and network-related software development, and deployment activities resulting in an operational environment prototype are expected to be part of the proposed activities.**
- **Proposals are encouraged to perform experimental deployment, protocol prototyping and testing, and evaluation using FABRIC (https://www.fabric-testbed.net).**

CC* Area#4 – Campus Computing and the Computing Continuum

- This program area promotes coordinated approaches in scientific computing at the campus level. This area invests in the seeding of new and shared computing resources at the campus level through investments in capacity computing in campus clusters. The program promotes a coordinated approach incentivizing multi-campus and national resource sharing.
- All proposals into this area must address:
  - Scientific and engineering projects and their research computing needs, describing project-specific scenarios for scientific computing tied to the proposed computing resources;
  - Features, capabilities, and software platforms representing the proposed computing resources; and
  - Scientific computing codes expected to run on the resources.
- NSF encourages proposals in this program area from under-resourced institutions and strong preference will be given to proposals demonstrating a compelling need for access to campus/cloud resources, including institutions lacking necessary computing and storage resources on campus.
- Proposals are required to commit to a minimum of 20% shared time on the cluster and describe their approach to making the cluster available as a shared resource external to the campus.

OSG - a National Fabric of Distributed Computing Services

https://osg-htc.org/campus-cyberinfrastructure

OSG About CC* News Events Contact Docs

349 Million Core Hours Contributed by CC* Campuses

In the last year, these campuses contributed over 349 million core hours to researchers using distributed high throughput computing (dHTC). (August 2020 - July 2021) These core hours supported over 236 projects in fields of study ranging from the medical sciences to the evolutionary sciences, and from biostatistics to physics. This campus support throughout the United States contributed to the advancement of science and to researchers both on and off their campuses. The campuses currently contributing CC*-funded resources are shown in the map below. Every month, the OSG is working to help additional campuses to join this effort to support open science by sharing their resources:

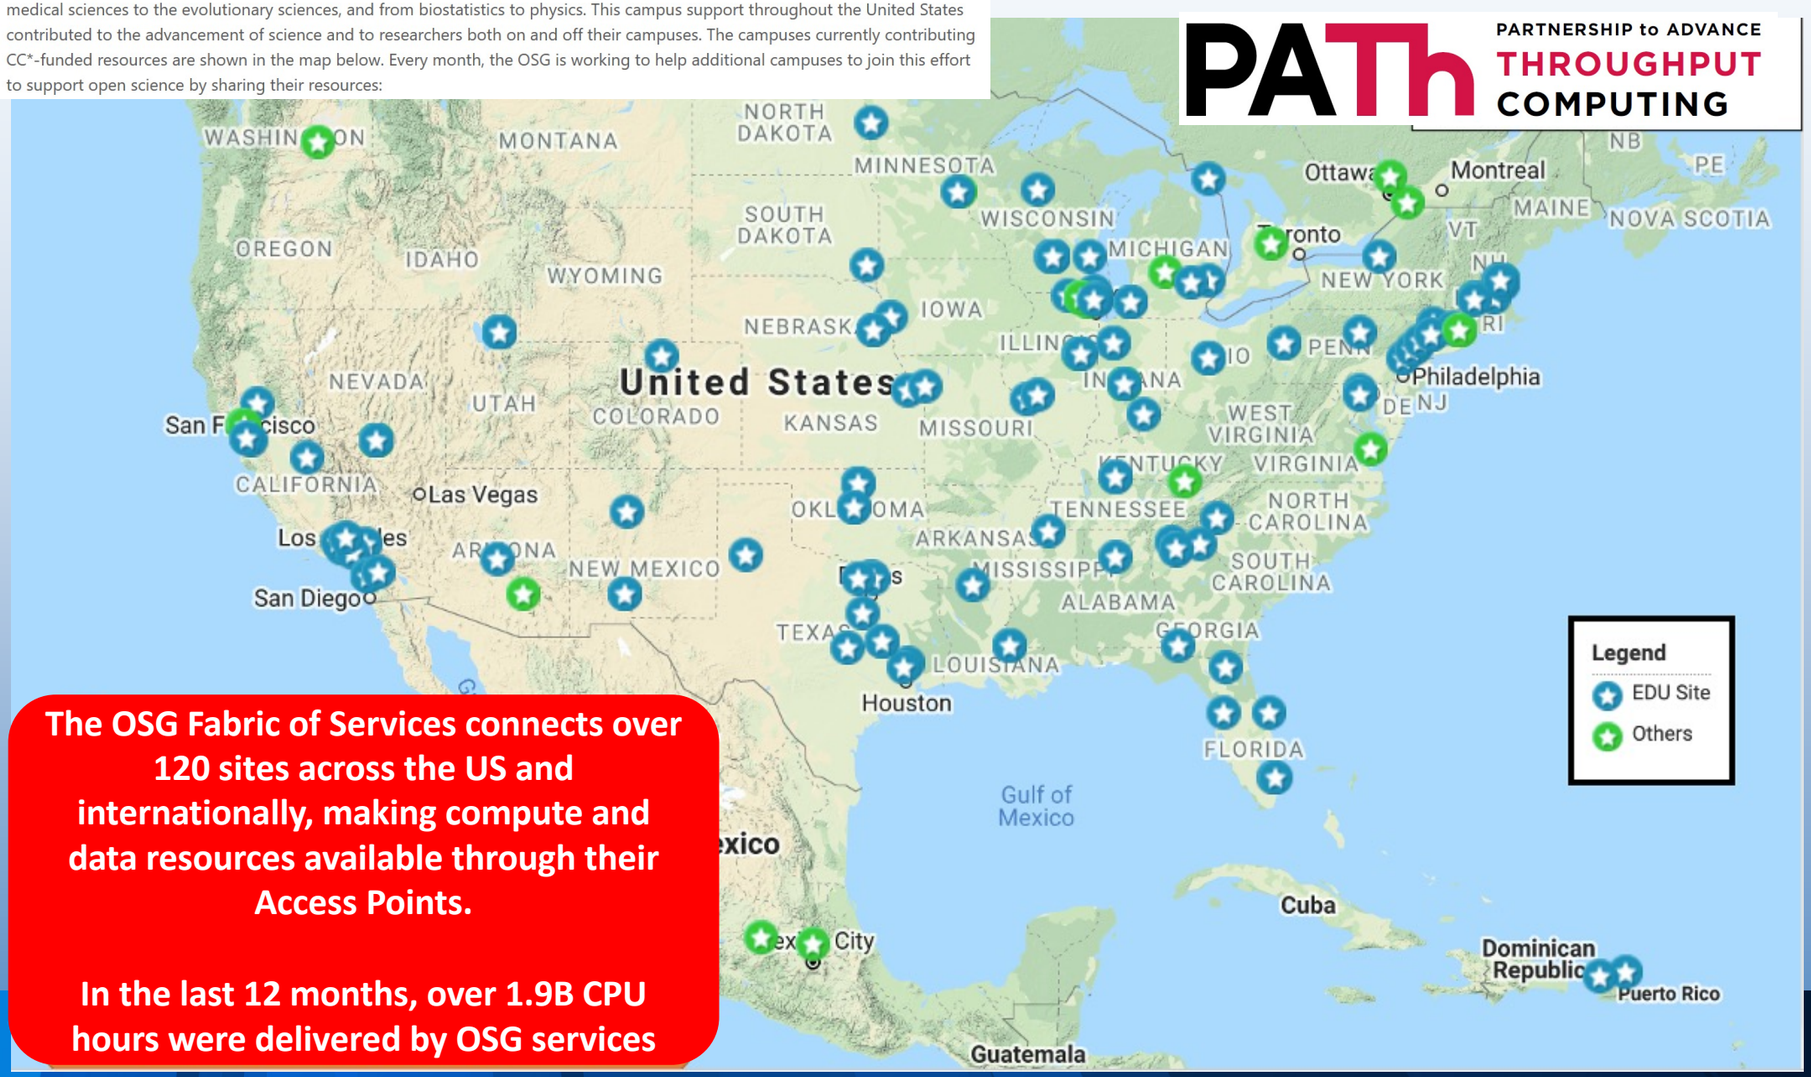

The OSG Fabric of Services connects over 120 sites across the US and internationally, making compute and data resources available through their Access Points.

In the last 12 months, over 1.9B CPU hours were delivered by OSG services

CC* Area#5 – Planning Grants and CI-Research Alignment (CIRA)

- This program area supports planning and coordination activities.
- Planning Grants:
  - Planning proposals should define a clear set of goals and a set of coordination and planning activities to meet those goals. Equipment costs are not allowed as part of a Planning Grant, and proposed costs are expected to include support for community coordination and planning activities
  - Planning proposals are welcome for areas (1) through (4) in CC*
- CI-Research Alignment (CIRA) proposals:
  - A CIRA award provides opportunities to foster new collaborations, including international partnerships, and address interdisciplinary topics.
  - A CIRA proposal is expected to develop a comprehensive CI strategy encompassing a campus, multiple campuses, or a state or regional research and education network entity.
  - The CIRA activity may encompass planning for a future CC* proposal, but goes beyond a specific campus network design, assessment of campus computing needs, or compilation of demanding science drivers to address integrated CI planning and scoping across the relevant scientific communities on campus, across multiple campuses, state-wide, or regionally.

Selected URLs from CC* 21-528

- http://fasterdata.es.net/campusCIplanning/
  - Example Campus CI Plans from CC* awardees
- http://fasterdata.es.net/science-dmz/
  - Reference material on re-architecting a campus network border
- https://epoc.global/cc/
  - NSF-funded project providing objective community expertise
- http://thequilt.net
  - See next talk by Jen Leasure
- http://www.opensciencegrid.org
  - OSG, a national distributed shared computing fabric

Thank you!

Full proposals due October 11, 2021, 5pm (submitter's local time)!!!
https://www.nsf.gov/pubs/2021/nsf21108/nsf21108.jsp?org=NSF

CONTACT INFORMATION

Kevin L. Thompson (kthompso@nsf.gov),
Deepankar Medhi (dmedhi@nsf.gov)
Program Directors, CC*

Subrata Acharya (acharyas@nsf.gov)
Timothy M. VanReken (tvanreke@nsf.gov),
Program Directors, EPSCoR